

## SEAMLESS ARBITRARY DATA INSERTION FOR STREAMING MEDIA

### *Technical Field*

The present invention relates to a methodology for inserting data such as  
5 advertising in streaming media and, more particularly, to a method for inserting such data  
at arbitrary locations along the stream.

### *Background of the Invention*

Traditional methods of delivery of multimedia data to end users fall into two  
10 categories: 1) broadcast industry methods and 2) computer industry methods. Broadcast  
methods (including motion picture, cable, television network, and record industries)  
generally provide storage in the form of analog or digitally recorded tape. The playing  
of tapes causes isochronous data streams to be generated which are then moved through  
broadcast industry equipment to the end user. Computer methods generally provide  
15 storage in the form of disks, or disks augmented with tape, and record data in compressed  
digital formats such as DVI, JPEG and MPEG. On request, computers deliver non-  
isochronous data streams to the end user, where hardware buffers and special application  
code smooths the data streams to enable continuous viewing or listening.

“Streaming media” is now considered to be a term of art and defines the  
20 capability to download multimedia files in real time (or “near real” time) and play the file  
on an end-user device, such as through a set-top box on a television or on a computer  
display device.

In the current state of the technology, it is possible to also download  
advertisements (or other media) to end users once the video or other presented material in  
25 the stream is completed. The need to “wait” until the end of the program thus limits the  
number, as well as the duration of advertisements that can be presented to end users.

Thus, a need remains in the art for an arrangement capable of inserting other  
video sources into a streaming media download.

### *Summary of the Invention*

The need remaining in the art is addressed by the present invention, which relates to a methodology for inserting data such as advertising in streaming media and, more particularly, to a method for inserting such data at arbitrary locations along the stream.

5 In accordance with the present invention, a “caching inserter” system is disposed at a local point of presence (POP) which intercepts a video stream destined for an end user. When insertion of advertising (or other information – such as a “news” bulletin weather information, for example) is desired, the video stream is stored/buffered at the local POP and the advertisement is played. When the ad is completed, the first stream is  
10 then played out of the buffer. Once the video has begun to enter and fill the buffer (that is, once an ad is inserted in the stream), the buffer must remain in the signal path for the duration of the video stream.

In an alternative embodiment of the present invention, the ad may first be loaded into a buffer at the POP and once the ad has begun, the concatenation point of the video  
15 stream can be appended to the end of the advertisement file. In this case, the length of the advertisement need not be monitored and no further action is required, since the video stream will simply re-commence when the ad is finished.

The POP where the advertisement cache is located could be at any convenient place along the signal path between the information source(s) and the end-user. For  
20 example, the advertisement cache could be at the headend, or at a FTTH node in a neighborhood. At the extreme, the advertisement cache could be part of a set-top box at the end-user location. Obviously, the location of the advertisement cache then affects the number of separate end users that can receive the same advertisement. Particularly, the  
25 further upstream the advertisement cache is located, the greater the broadcast potential of a particular ad. Advantageously, however, judicious location and control of various advertisement contents can be incorporated into the present invention such that content can be targeted to particular end users.

It is an advantage of a particular arrangement of the present invention that the requirement of passing an advertisement stream through a cache allows for various  
30 statistics to be collected at the POP, such as the number of times a particular ad is played, the number of end users that viewed the ad, the time of day and day of week the ad was

played, etc. Such information can then be collected by a billing system and used to derive price information for use with the advertisement suppliers.

Various other features and advantages of the present invention will become apparent during the course of the following discussion and by reference to the accompanying drawings.

### ***Brief Description of the Drawings***

Referring now to the drawings,

FIG. 1 is a system-level block diagram illustrating the capability of inserting advertisement into streaming multimedia in accordance with the present invention;

FIG. 2 contains two separate diagrams of media stream, one including advertisement as a prefix and the other including advertisement embedded within the multimedia stream; and

FIG. 3 illustrates an exemplary cache/inserter arrangement that may be utilized to implement the arrangement of the present invention.

### ***Detailed Description***

An exemplary architecture for supporting the insertion of an alternative message into streaming multimedia in accordance with the present invention is illustrated in FIG.

1. For the purposes of discussion, the “alternative message” will be referred to as an “advertisement” and, in one embodiment, may be advertisements that are targeted to the interests of the specific subscriber/end-user (hereinafter referred to simply as an “end-user”), where the phrase “end-user” is considered to include both a single end-user and a group of end-users (as in the case of a “multicast”/“broadcast”). However, it is to be understood that various other types of messages may also be inserted into a streaming multimedia download. For example, news or weather emergency messages may be inserted, if warranted, into an on-going multimedia stream.

Referring to FIG. 1, an end-user 10 is depicted as a residential end-user that is capable of receiving multimedia information through either a television 12 (using a set-top box 14 or other appropriate residential gateway device) or a computer terminal 16 (including appropriate software to convert the received files and play them on the display

of computer 16). In a preferred architecture, a coaxial cable connection 18 is used to provide communication between end-user 10 and a head-end, or local distribution network 20 (alternatively, connection 18 may comprise a DSL connection or a wireless connection). Although not depicted in the diagram of FIG. 1, local distribution network 20 is similarly connected to a multiplicity of end-users and, in fact, may be connected to these users in a tree-and-branch architecture. For conventional communications, local distribution network 20 is connected through one or more links 22 (in most cases, optical fiber links) to a data network 24 (which may be referred to as the Internet). Multiple sources of information, resident on servers, are connected through their own links to data network 24, and other sources are likewise interconnected and are able to communicate with each other using data network 24. For the purposes of the present discussion, one such multimedia server 26 and a communication link 28 are shown in FIG. 1 as connected to data network 24, where the content stored on multimedia server 26 is thus accessible by end-user 10. In a conventional use of this arrangement, end-user 10 would send a request through local distribution network 20 and into data network 24 for one or more multimedia files to be downloaded and “played” on one of his devices (for example, television 12 or computer 16). A link would be established between end-user 10 and the appropriate multimedia server 26 holding the file requested by the end-user. With the link established, the requested file is “streamed” down the communication links to end-user 10 in a continuous fashion, allowing for uninterrupted viewing (of, for example, a movie, music video, or the like) of the multimedia content of the file.

In accordance with the present invention, additional components have been implemented in the communication path to allow for advertisements (or other messages, as mentioned above) to be inserted in streaming multimedia in a seamless fashion, that is, without the need to “wait” for predetermined time slots or breaks in the multimedia stream. Referring to FIG. 1, a “caching inserter” 30, located in a local point-of-presence (POP) 32 along the path between data network 24 and local distribution network 20 functions to monitor incoming streaming multimedia along link 34 and insert, as discussed in detail below, an advertisement, at a location along the stream. The advertisement medium itself may be resident on a separate advertisement content server 36 coupled by a separate data link 38 to data network 24. A “rule” server 40 is shown in

FIG. 1 and may be used to direct particularly advertising content to specific end-users, where rule server 40 interacts with caching insert 30, as shown in FIG. 1. In order to control payment for the advertisement insertion service, a separate billing system 42 may also be associated with caching inserter 30.

5 In operation, a specific end-user first requests a particular multimedia file to be accessed and streamed into one of his devices. Referring to FIG. 1, such a request first passes through caching inserter 30 and then enters data network 24, which seeks out a specific server 26 holding the desired multimedia file. In order to identify a flow that contains streaming multimedia content, caching server 26 may detect a request or,  
10 through examination of the content of the flow, detect the download of a multimedia file. Once caching inserter 30 receives the initial request, the identity of end-user 10 may be forwarded to rule server 40 that can then access its internal database to determine the type of advertisement (or other information) that would most likely be of interest to that user. For example, end-user 10 may be interested in antique cars and that information would be  
15 stored in “rule” server 40, along with the IP address(es) of servers holding advertisement information related to antique cars. Rule server 40 then sends the IP address information back to caching inserter 30 which then accesses data network 34 to retrieve a particular advertisement multimedia file (from, for example, advertising content server 36). The retrieved advertisement file is then forwarded to and cached at caching inserter 30.

20 In accordance with the present invention, once caching inserter 30 has an advertisement file available (either stored, or available as a stream from the advertisement server), the incoming (requested) multimedia file destined for end-user 10 can be cached, and the advertisement multimedia file streamed in its place to end-user 10 in a seamless manner. The multimedia file itself is not stopped at any point, but is  
25 continuously cached by caching inserter 30 as it is received from streaming media server 26. At the end of the advertisement file, the stored multimedia stream is played out and, for the duration of the stream, will continue to pass through a cache, since there will continue to be a time delay between the incoming stream to caching inserter 30 and the outgoing stream to end-user 10 (the “delay” being the length of the inserted  
30 advertisement). Since the end time of the advertisement is known by caching inserter 30, the streaming multimedia file can merely be appended to the end of the advertisement,

allowing the playing of the requested file to recommence in a seamless manner without further action required by caching inserter 30.

As shown in FIG. 1, a billing system 42 may be included in the arrangement and configured to receive input information from caching inserter 30. That is, billing system 42 may be used to collect information regarding the identity of the various advertisement files that are requested and the number of times each file is requested. Other information, such as the time of day of the request, day of the week, etc., may also be collected. This information is useful both for the supplier of the advertisements and for billing the advertisement for the supplied service (e.g., the identity of the end-user(s) viewing the advertisement, the number of times a particular ad is requested in a day/week, etc.).

Although only one such end-user is shown in FIG. 1, it is to be understood that caching inserter 30 of the present invention is to be used with multiple users, where rule server 40 includes in a database information of a plurality of known end-users and is capable of requesting various types of files that are pertinent to different end-users. Moreover, caching inserter 30 itself is capable of caching any desired number of separate advertisements and controlling the insertion of the advertisements in the proper streaming multimedia files, while at the same time caching and then controlling the sending out of the cached multimedia files to the appropriate end-users.

In one embodiment of the present invention, the multimedia stream will begin to reach the end-user before an advertisement is inserted, as represented by diagram A in FIG. 2. In this case, an initial header slot 50 is understood by the receiving device (set-top box or computer) as the requested multimedia file. The receiving device is then properly configured to play the streaming multimedia file, which follows as segment 52 of diagram A. At some point along the playing of the file, an advertisement is inserted as segment 54 and the remainder of the multimedia stream, represented as segment 56, is appended to the end of the advertisement, as discussed above. In one example, a timing signal in the multimedia file may be used as a "trigger" to initiate the insertion of the advertisement into the stream. Alternatively, the advertisement may simply be cached and held until the end of the multimedia stream, as shown in diagram B of FIG. 2, or played as the initial segment (between the header and the streaming multimedia), as shown in diagram C. In any alternative, additional bytes may be added to the beginning and end of

the advertisement file to create a “fade in” and “fade out” between the advertisement and requested multimedia stream, where such fading is considered to be more acceptable by the end-user than a sharp break between the independent streams.

FIG. 3 illustrates an exemplary caching inserter 30 that may be used in the system of the present invention as described above in association with FIG. 1. Included within caching inserter 30 is a first cache 60 used to store an incoming specific advertisement file, designated as stream A, from advertisement content server 60. A second cache 62 is used to store the incoming multimedia stream, designated as stream MM, where as discussed above, cache 62 is used to store incoming stream MM once advertisement stream A has begun to be “played out” toward end-user 10. An analyzer 64 is included in the communication path between the network (which may be data network 24, local distribution network 20, or any other communication network, since caching inserter 30 may be disposed at any appropriate location along the path between the various multimedia servers and end-users) and end-user 10. Analyzer 64 will detect when a multimedia stream has been received, determine the associated protocol and rate, and transmit this information to a timing and control unit 66. Other control inputs to timing and control unit 66 are from advertising content (first) cache 60 and multimedia content (second) cache 62, where control unit 66 uses the input from these caches to determine the proper stream to be played out to end-user 10. For example, unless and until a full advertisement file is loaded in first cache 60, multimedia stream MM can be played coupled in real time to end-user 10, indicated as position (1) of switch 67 of inserter 30. When control unit 66 receives a particular indication from first cache 60 (such as “EOF” received), it knows that advertisement stream A is ready to be played out and multimedia stream MM can be held. A signal is then sent from timing and control unit 66 to switch 67, requesting the switch to move to position (2), allowing the cached advertisement stream A to be sent to end-user 10. Once this process begins, multimedia stream MM will begin to be stored in second cache 62. When timing and control unit 66 receives a command that the end of the advertisement has been reached, a signal to move switch 67 to position (3) will be sent, allowing the cached multimedia stream to be sent to end-user 10. If the advertisement is inserted at the beginning of or during the multimedia stream, the received multimedia stream will need to be stored in, then played out of, second

cache 62 for the duration of the multimedia file. A rate and/or format transcoder 68 may also be included in the signal path between first cache 60 and end-user 10 to insure that the advertisement is transmitted at the same data rate and format as the requested multimedia file, where this information has previously been retrieved by analyzer 64, or  
5 by some other means.

While the arrangement as described above illustrates the various servers as being in geographically disparate locations, it is to be understood that one or more may be co-located. Additionally, as noted above, the caching inserter of the present invention may be disposed at any location between the various content servers and the end-user. Indeed,  
10 various other modifications to the arrangement as described above will occur to those skilled in the art and are considered to fall within the spirit and scope of the present invention as defined by the claims appended hereto.